arXiv:1701.02512v1 [stat.ME] 10 Jan 2017

An RKHS model for variable selection in functional regression

José R. Berrendero, Beatriz Bueno-Larraz, Antonio Cuevas
Departamento de Matemáticas
Universidad Autónoma de Madrid, Spain

Abstract

A mathematical model for variable selection in functional regression models with scalar response is proposed. By "variable selection" we mean a procedure to replace the whole trajectories of the functional explanatory variables with their values at a finite number of carefully selected instants (or "impact points"). The basic idea of our approach is to use the Reproducing Kernel Hilbert Space (RKHS) associated with the underlying process, instead of the more usual $L^2[0,1]$ space, in the definition of the linear model. This turns out to be especially suitable for variable selection purposes, since the finite-dimensional linear model based on the selected "impact points" can be seen as a particular case of the RKHS-based linear functional model. In this framework, we address the consistent estimation of the optimal design of impact points and we check, via simulations and real data examples, the performance of the proposed method.

Keywords: feature selection, functional regression, impact points, variable selection.

1 Introduction: statement of the problem and motivation

The problem under study: variable selection in functional regression

The study of regression models is clearly among the leading topics in statistics. In particular, these models play a central role in the theory of statistics with functional data, often called Functional Data Analysis (FDA); see Cuevas (2014) for an overview on FDA.

Throughout this paper, we will consider "functional data" consisting of independent $X_1 = X_1(t), \ldots, X_n = X_n(t)$ observations (trajectories) drawn from a second-order (L^2) stochastic process $X = X(t)$, $t \in [0,1]$, with continuous trajectories, continuous mean function $m = m(t)$ and covariance function $K(s,t)$. All the involved random variables are supposed to be defined in a common probability space $(\Omega, \mathcal{A}, \mathbb{P})$.

We are interested on regression models with scalar response, of type $Y_i = g(X_i) + \varepsilon_i$, where g is a real function defined on a suitable space $\mathcal{X}$ where the trajectories of our process are supposed to live and ε_i are independent errors (and also independent from the X_i) with mean zero and common variance σ^2.

More specifically, we are concerned with variable selection issues; see, Fan and Lv (2010), Berrendero et al. (2016, Sec. 1) for additional information and references. Basically, a variable selection functional method is an automatic procedure that takes a function $\{x(t),\ t \in [0,1]\}$ to a finite-dimensional vector $(x(t_1), \ldots, x(t_p))$. The overall idea for variable selection is to choose the variables $x(t_i^*)$ (or, equivalently, the"impact

points" $t_1, \ldots, t_p \in [0,1]$; see Kneip et al. (2016)), in an "optimal way" so that the original functional problem (regression, classification, clustering,...) is replaced with the corresponding multivariate version of this problem, based on the selected variables. In the regression problem, this would amount to replace the functional model $Y_i = g(X_i)+\varepsilon_i$ by a finite dimensional version of type $Y_i = \phi(X(t_1), \ldots, X(t_p)) + e_i$. Nevertheless, note that still the problem is of a functional nature, since the variable selection process of the t_i involves in principle the full data trajectories.

Some motivation. The classical linear L^2-model. Its drawbacks for variable selection purposes

It is quite natural to assume that the explanatory functional variables $X_i = X_i(t)$ are members of the space $L^2[0,1]$, endowed with the usual inner product $\langle x_1, x_2\rangle_2 = \int_0^1 x_1(t)x_2(t)dt$, for $x_1, x_2 \in L^2[0,1]$. In this setting, the most popular choice for g is, by far, a linear (or affine) operator from $L^2[0,1]$ to $\mathbb{R}$ which leads to a model of type

$$Y_i = \alpha_0 + \langle X, \beta\rangle_2 + \varepsilon_i,\ i = 1, \ldots, n, \tag{1}$$

where $X = X(t)$ is the explanatory functional variable, $\alpha_0 \in \mathbb{R}$ is the intercept constant and $\beta \in L^2[0,1]$ denotes the slope function. As in the standard multivariante regression model, the aim here is to estimate α_0 and β in order to be able to make accurate predictions of the response variable Y.

The corresponding theory is outlined in several places; see, e.g., the article Cardot and Sarda (2010) or the books Ferraty and Vieu (2006) and Horváth and Kokoszka (2012). The Hilbert structure of the $L^2[0,1]$ space allows us to keep ourselves as close as possible to the usual least square framework in multivariate regression; for example, the projection $P(x)$ of an element x on a subspace H is characterized by the orthogonality condition $\langle x - P(x), a\rangle_2 = 0$, for all $a \in H$. However, other crucial differences with the finite-dimensional case (mostly associated with the non-invertibility of the covariance operator of the process $X(t)$) makes the functional L^2 regression theory far from trivial. Most of these difficulties are intrinsic to the infinite-dimensional nature of the data, so that they cannot be overcome by just replacing $L^2[0,1]$ with another function space. However, when it comes to variable selection applied to linear regression, it would be useful to have the finite dimensional linear model (based on the selected variables)

$$Y_i = \alpha_0 + \sum_{j=1}^{p} \beta_j X_i(t_j) + \varepsilon_i,\ i = 1, \ldots, n \tag{2}$$

as a particular case of our general model. Notice that (1) cannot be established in the L^2 framework, since a transformation of type $x \in L^2[0,1] \mapsto \sum_{i=1}^{p} \beta_i x(t_j)$ is not a linear continuous functional in L^2. In heuristic terms, one would need to look for the regression function β in a suitable space, for which a "finite-dimensional" model such as (2) could make sense. More precisely, we will replace the $L^2[0,1]$ space in which β is assumed to live with the Reproducing Kernel Hilbert Space (RKHS), $\mathcal{H}(K)$, associated with K. As we will comment in Section 2.2 this entails some technical difficulties, since the sample space for the data $X = X(t)$ is still $L^2[0,1]$.

The assumed membership of β to $\mathcal{H}(K)$ entails some additional restrictions of regularity on the slope function β (when compared to the simple assumption $\beta \in L^2[0,1]$). In any case, some restrictions on β appear also in different ways when the classical L^2-model (1) is considered. The reason is that the space $L^2[0,1]$ is in fact too large from several points of view. Hence, in spite of the advantages commented above, one typically uses penalization or projection methods to exclude extremely rough solutions in the estimation of β.

Our proposal here, as presented in the next section, aims at reconciling two targets: first, we look for a functional linear model, wide enough to include finite-dimensional versions, such as (2), as particular cases. Second, we would like to achieve such goal with a minimal change in the "parameter space" where β lives.

Some notation

A set of possible "impact" points $t_1, \ldots, t_p \in [0,1]$ will be denoted T (sometimes S) or T_p when we want to stress the cardinality of T. Also, $X(T_p)$ will stand for $(X(t_1), \ldots, X(t_p))'$. The superindex * will be used to denote that the points t_i^* are the "true" ones, or the "optimal" ones according to some criterion.

Given a random variable Z (with finite variance) the notation Z_{T_p} will refer to the L^2-projection of Z on the space spanned by the components of $X(T_p) - m(T_p)$.

If $p^* < p$, the notation $T_{p^*} \prec T_p$ will indicate that all the points in T_{p^*} belong also to T_p.

Finally, as usual in statistics, we use a hat to denote the estimated quantities (or the predicted variables). For instance, $\widehat{T}_p$ will denote a data-driven estimator of T_p and $\widehat{Y}_{\widehat{T}_p}$ will stand by the corresponding (fully data-driven) prediction of the response Y_{T_p}. The halfway notation $Y_{\widehat{T}_p}$ will represent the projection of the response variable onto the space spanned by the marginal variables indexed by the estimated points $\widehat{T}_p$.

Organization of the paper

In Section 2 we introduce and motivate (in population terms) our variable selection procedure. The asymptotic properties of the empirical version (when the parameters are estimated) are considered in Section 3. The problems associated with the choice of the number p of selected variables are analyzed in Sections 4 and 5. The empirical results (simulations and real data examples) are presented in Section 6. Some technical proofs are included in Section 8.

2 An RKHS-based linear model suitable for variable selection

Our choice of the ambient space for the slope function β is, in some sense, "customized" for the problem at hand, since we will consider the Reproducing Kernel Hilbert Space (RKHS) associated with the process $\{X(t),\ t \in [0,1]\}$.

The theory of RKHS goes back to the 1950's; see Janson (1997, Appendix F) for details and references. It has found a surprisingly large number of applications in different fields, including statistics, see Berlinet and Thomas-Agnan (2004).

2.1 RKHS spaces in a nutshell

Before establishing our RKHS-based regression model we need a minimal background on RKHS. Our starting point will be our underlying L^2-process $X = \{X(t)\, t \in [0,1]\}$ with a continuous strictly positively definite covariance function $K(s,t)$ and a continuous mean function $m(t)$.

We first introduce an auxiliary space, associated with K, which we will denote by $\mathcal{H}_0(K)$. It is defined by the set of all finite linear combinations of type $\sum_i^n a_i K(s,t_i)$, that is,

$$\mathcal{H}_0(K) := \{f :\ f(s) = \sum_{i=1}^{n} a_i K(s,t_i),\ a_i \in \mathbb{R},\ t_i \in [0,1],\ n \in \mathbb{N}\}.$$

In such space we define an inner product $\langle \cdot, \cdot \rangle_K$ by $\langle f, g \rangle_K = \sum_{i,j} \alpha_i \beta_j K(s_j, t_i)$, where $f(x) = \sum_i \alpha_i K(x, t_i)$ and $g(x) = \sum_j \beta_j K(x, s_j)$.

Now, the RKHS associated with K, denoted by $\mathcal{H}(K)$, is defined as the completion of $\mathcal{H}_0(K)$. More precisely, $\mathcal{H}(K)$ is the set of functions $f : [0,1] \to \mathbb{R}$ obtained as t-pointwise limits of Cauchy sequences $\{f_n\}$ in $\mathcal{H}_0(K)$; see Berlinet and Thomas-Agnan (2004, p. 18).

As a conclusion, in heuristic terms, one could say that $\mathcal{H}(K)$ is made of all linear combinations of type $f(s) = \sum_i^n a_i K(s,t_i)$ plus all the functions which can be obtained as limits of them. A natural question is when we can ensure that we have identifiability in this space. It is easy to see that the elements of $\mathcal{H}_0(K)$ have a unique representation in terms of K whenever K is strictly positive definite. For additional details we refer again to Janson (1997, Appendix F).

Among the many interesting properties of RKHS spaces, let us especially recall two which will be particularly useful in what follows.

Reproducing property. $f(t) = \langle f, K(\cdot, t) \rangle_K$, for all $f \in \mathcal{H}(K)$, $t \in [0,1]$.

Natural congruence. Denote by $\mathcal{L}_X$, the linear (centered) span of X (i.e. the family of finite linear combinations of type $\sum_i \lambda_i (X(t_i) - m(t_i))$). Let $\overline{\mathcal{L}}_X$ be the L^2-completion of $\mathcal{L}_X$. It is clear that $\overline{\mathcal{L}}_X$ is a closed subspace of the usual Hilbert space $L^2(\Omega)$ of random variables with a finite second moment; this can be seen as the minimal Hilbert space including the variables $X(t)$. It can be proved, see Berlinet and Thomas-Agnan (2004, Th. 35), that $\Psi_X(\sum_i a_i (X(t_i) - m(t_i)) = \sum_i a_i K(\cdot, t_i)$ defines (when extended by continuity) a congruence between $\overline{\mathcal{L}}_X$ and $\mathcal{H}(K)$. This means that the extension of Ψ_X is a linear bijective transformation which preserves the inner product. Such congruence is often called *Loève's isometry*. In explicit terms (see Lukić and Beder (2001, Lemma 1.1)), Loève's isometry between $\overline{\mathcal{L}}_X$ and $\mathcal{H}(K)$ can be defined by

$$Y \mapsto \Psi_X(Y) = \Psi_X(Y)(t) = \mathbb{E}(Y(X(t) - m(t))) \tag{3}$$

So, the RKHS space $\mathcal{H}(K)$ is an isometric copy of $\overline{\mathcal{L}}_X$. In this sense, it contains the "Dirac deltas" $X(t_i) - m(t_i)$ at t_i (as $\mathcal{H}(K)$ contains $K(\cdot, t_i)$). Both spaces $\overline{\mathcal{L}}_X$ and $\mathcal{H}(K)$ can be identified.

There is, however, a not-so-nice feature in the RKHS space associated with the process $X(t)$: under very general conditions, this space does not contain, with probability one, the trajectories of the process X; see, e.g., Pillai et al. (2007, Th. 11), (Lukić and Beder, 2001, Cor. 7.1). This will have some consequences in the formulation of our regression model, as pointed out below.

2.2 The RKHS functional regression model

We propose to replace the standard L^2 functional regression model (1) with the following RKHS counterpart

$$Y_i = \alpha_0 + \langle X_i, \beta \rangle_K + \varepsilon_i, \ i = 1, \ldots, n \tag{4}$$

where $\beta \in \mathcal{H}(K)$ and $\langle \cdot, \cdot \rangle_K$ denotes the inner product in $\mathcal{H}(K)$.

Since the estimation of the intercept term α_0 is straightforward from those of β and m, we will assume, without loss of generality, that $\alpha_0 = 0$ in what follows.

As mentioned at the end of the previous subsection, it is important to keep in mind that the trajectories X_i of the process X do not belong to $\mathcal{H}(K)$. Thus, the expression $\langle X_i, \beta \rangle_K$ has no direct meaning, unless it is appropriately interpreted: in what follows, $\langle X_i, \beta \rangle_K$ must be understood as $\Psi_{X_i}^{-1}(\beta)$, where Ψ_X is defined in (3). Such an interpretation of $\langle X, \beta \rangle_K$ arises in the classical paper by Parzen (1961, Th. 7A), aiming at different statistical purposes. This situation is conceptually similar to that arising in the definition of Itô's integral, so that $\langle X, \beta \rangle_K$ can be interpreted as a stochastic integral. In fact, when $X(t)$ is a standard Brownian Motion model (4) is equivalent to

$$Y = \int_0^1 \beta(s) dX(s) + \varepsilon, \text{ with } \beta \in \mathcal{H}(K),$$

where $\int_0^1 \beta(s) dX(s)$ is Itô's integral.

2.3 Variable selection in the RKHS functional regression model

Consider the RKHS functional regression model (4) introduced in the previous paragraph, where $\mathbb{E}(\varepsilon) = \mathbb{E}((X(t) - m(t))\varepsilon) = 0$ and $\mathrm{Var}(\varepsilon) = \sigma^2$.

Our goal

Under this model, for fixed p, we aim at selecting p values $t_1, \ldots, t_p$ in order to use the p dimensional vector $(X(t_1), \ldots, X(t_p))$ instead of the whole trajectory $\{X(t) : t \in [0, 1]\}$ in our regression problem. Formally, we want to establish a transformation

$$\{X(t) : t \in [0, 1]\} \mapsto (X(t_1), \ldots, X(t_p)),$$

which should be "optimal" in the sense that the points $(t_1, \ldots, t_p)$ are chosen according to an optimality criterion, oriented to minimize the information loss in the passage from infinite to finite dimension.

In this section, we address this problem at the population level, that is, we assume that the parameters defining the model (the regression function β, the covariance function K of the process X, the mean function m and the variance of the error variable, σ^2) are known. Of course, the practical implementation will require using suitable estimators of the unknown parameters. This raises several questions concerning the sample behavior of the method which will be addressed in subsequent sections.

The optimality criterion Q_1

The first obvious question to address in such strategy is the choice of the optimality criterion. We will see that, in fact, different criteria can be used but, fortunately, they are all equivalent.

One of the basic goals of a functional regression model is to predict the value of the response variable for a given realization of the regressor function. Then, a sensible approach for variable selection is to choose the p points $(X(t_1), \ldots, X(t_p))$ that give the best linear prediction (in the sense of the L^2 norm) of Y. This implies to find the vector T_p that minimizes the function

$$Q_1(T_p) := \min_{(\beta_1,\ldots,\beta_p)\in\mathbb{R}^p} \|Y - \sum_{j=1}^{p} \beta_j(X(t_j) - m(t_j))\|_2^2. \tag{5}$$

Where to look for the optimum

An important technical aspect is the choice of an appropriate subset $\Theta_p \subset [0,1]^p$ to look for the optimum of the continuous function Q_1. This subset must be compact in order to guarantee the existence of the optimum. Moreover, if we want to get a meaningful optimal value of $T_p = (t_1, \ldots, t_p)$ we should rule out those points including repeated values in the coordinates t_i. To this end, we will fix an arbitrarily small fixed value $\delta > 0$, and will look for our optimum in the space

$$\Theta_p = \Theta_p(\delta) = \{T_p = (t_1, \ldots, t_p) \in [0,1]^p : \, t_{i+1} - t_i \geq \delta, \text{ for } i = 0, \ldots, p\}, \tag{6}$$

where $t_0 = 0$, $t_{p+1} = 1$. In practice, the restriction to the subset Θ_p is not relevant, since we observe the functions in a finite grid, and we can set $\delta > 0$ as small as required so that all the points in the grid belong to Θ_p.

The reason for the choice (6) of Θ_p is technical, very much in the same spirit of Ji and Müller (2016, Eq. (9)). We need to work on a compact set and, at the same time, to avoid degeneracy problems in the choice of the points $(t_1, \ldots, t_p)$ that could lead to a singular covariance matrix in $(X(t_1), \ldots, X(t_p))$. Other choices are possible for Θ_p. For example, one could think of defining $t_0 = 0$, $t_{p+1} = 1$ and

$$\Theta_p = \{T_p = (t_1, \ldots, t_p) \in [0,1]^p : \, t_i \leq t_{i+1}, \text{ for } i = 0, \ldots, p\} \tag{7}$$

This could lead to "degenerate" options with $t_i = t_{i+1}$ for some values of i. However, the theory we develop below using (6) could be carried out alternatively with (7) as long as we adopt the criterion of reducing the dimension of those vectors $(t_1, \ldots, t_d)$ with ties in the coordinate values by keeping just one coordinate for each different value. In this way, for example, $(1/3, 1/3, 1/2, 3/4)$ would be interpreted just as $(1/3, 1/2, 3/4)$.

Two additional, equivalent optimality criteria

A second optimality criterion, equivalent to that based on Q_1, arises if we take into account that, from the reproducing property, when the regression function is a finite linear combination of the form $\sum_{i=1}^{p} \beta_i K(t_i, \cdot)$, model (4) reduces to the usual finite dimensional multiple regression model:

$$Y = \sum_{i=1}^{p} \beta_j (X(t_j) - m(t_j)) + \varepsilon. \tag{8}$$

Then, another sensible approach for variable selection is to choose those points $t_1, \ldots, t_p$ giving the best approximation of the true regression function β in terms of a finite linear combination of the form $\sum_{i=1}^{p} \beta_i K(t_i, \cdot)$. It is quite natural to use the norm in $\mathcal{H}(K)$ to assess this approximation since both β and these finite linear combinations live in this RKHS. This approach amounts to find the vector $T_p \in \Theta_p$ that minimizes the function

$$Q_2(T_p) := \min_{(\beta_1, \ldots, \beta_p) \in \mathbb{R}^p} \|\beta - \sum_{j=1}^{p} \beta_j K(t_j, \cdot)\|_K^2. \tag{9}$$

Proposition 1 below shows that the variable selection procedures defined by (5) and (9), although apparently different, are indeed equivalent. Moreover, in the proof of Proposition 1 we will see that the minimum in the expressions of Q_1 and Q_2 is achieved at the value

$$(\beta_1^*, \ldots, \beta_p^*) = \Sigma_T^{-1} c_T,$$

where $c_T = (\mathrm{cov}(X(t_1), \ldots, \mathrm{cov}(X(t_p))'$ and Σ_T is the covariance matrix of X_T, for $T = (t_1, \ldots, t_p)$.

In addition, we show that the Q_1 and Q_2-based criteria are also both equivalent to a third criterion, defined in terms of a functional Q_0, which only depends on the covariances $K(t_i, t_j)$ and $\mathrm{Cov}(X(t_i), Y)$ for $i, j = 1, \ldots, p$. This Q_0 criterion turns out to be especially useful to implement the method in practice.

Proposition 1. *Assume that Y and X fulfil the RKHS functional regression model* (4)*. Then,*

$$\operatorname*{argmin}_{T_p \in \Theta_p} Q_1(T_p) = \operatorname*{argmin}_{T_p \in \Theta_p} Q_2(T_p) = \operatorname*{argmax}_{T_p \in \Theta_p} Q_0(T_p), \tag{10}$$

where Q_1 and Q_2 are defined in (5) *and* (9) *respectively, and*

$$Q_0(T_p) := c'_{T_p} \Sigma_{T_p}^{-1} c_{T_p}, \tag{11}$$

with $c_{T_p} = (Cov(X(t_1), Y), \ldots, Cov(X(t_p), Y))'$ *and* Σ_{T_p} *the* $p \times p$ *matrix with entries* $K(t_i, t_j)$.

Proof. Since $\mathbb{E}(\varepsilon) = 0$ and $\langle X(t), \epsilon \rangle_2 = 0$,

$$\|Y - \sum_{j=1}^{p} \alpha_j (X(t_j) - m(t_j))\|_2^2 = \|\langle X, \beta \rangle_K - \sum_{j=1}^{p} \alpha_j (X(t_j) - m(t_j))\|_2^2 + \sigma^2.$$

On the other hand, Loève's isometry implies

$$\|\langle X, \beta \rangle_K - \sum_{j=1}^{p} \alpha_j (X(t_j) - m(t_j))\|_2^2 = \|\beta - \sum_{j=1}^{p} \alpha_j K(t_j, \cdot)\|_K^2.$$

From the last two equations, it follows that $Q_1(T_p) = Q_2(T_p) + \sigma^2$ and hence the first equality in (10).

By the reproducing property,

$$\|\beta - \sum_{j=1}^{p} \alpha_j K(t_j, \cdot)\|_K^2 = \|\beta\|_K^2 - 2\sum_{j=1}^{p} \alpha_j \beta(t_j) + \sum_{i=1}^{p}\sum_{j=1}^{p} \alpha_i \alpha_j K(t_i, t_j). \tag{12}$$

The function K is positive semidefinite so that the last function is convex in $\alpha = (\alpha_1, \ldots, \alpha_p)$. By computing its gradient (with respect to α) it is very easy to see that the minimum is achieved at $\alpha^* = (\alpha_1^*, \ldots, \alpha_p^*) = \Sigma_{T_p}^{-1} \beta_{T_p}$. Then,

$$Q_2(T_p) = \|\beta - \sum_{j=1}^{p} \alpha_j^* K(t_j, \cdot)\|_K^2 = \|\beta\|_K^2 - \beta'_{T_p} \Sigma_{T_p}^{-1} \beta_{T_p}, \tag{13}$$

where $\beta_{T_p} = (\beta(t_1), \ldots, \beta(t_p))'$. Finally, using $\mathbb{E}(X(t)\varepsilon) = 0$ and Equation (3) we get

$$\mathrm{Cov}(Y, X(t)) = \mathbb{E}\big[\langle X, \beta \rangle_K (X(t) - m(t))\big] = \Psi_X(\langle X, \beta \rangle_K)(t) = \beta(t).$$

To obtain the last equality, recall that $\langle X, \beta \rangle_K = \Psi_X^{-1}(\beta)$. Therefore $\beta_{T_p} = c_{T_p}$ and, by (13), $Q_2(T_p) = \|\beta\|_K^2 - Q_0(T_p)$. This implies the second equality in (10). □

The criterion provided by Q_0 (or Q_1, Q_2) for variable selection was already considered by McKeague and Sen (2010), for $p = 1$, when $X(t)$ is a fractional Brownian Motion with Hurst exponent $H \in (0, 1)$, and by Ji and Müller (2016) for $p \geq 1$ in the usual L^2 functional regression model. The RKHS formalism we incorporate here provides a simple way to describe the scenario under which variable election would lead to the optimal solution (with no loss of information). Variable selection is specially suitable when the true regression model is sparse, meaning that the response depend on the explanatory variables through their values at a finite small number of p^* points. As it was mentioned before, this is the case under (4) when

$$\beta(t) = \sum_{j=1}^{p^*} \beta_j K(t_j^*, t). \tag{14}$$

Let $T^*_{p^*} = (t^*_1, \ldots, t^*_{p^*}) \in \Theta_{p^*}$. Then, it is clear that, under (14),

$$Q_2(T^*_{p^*}) = 0 \leq Q_2(T_{p^*}), \quad \text{for all } T_{p^*} \in \Theta_{p^*}.$$

As a consequence, the true set of relevant variables $T^*_{p^*}$ is the one selected by the optimization of the functions in Proposition 1. In this reasoning we have considered the case when we know the actual number of points p^* to be selected. In practice, this is not usually the case. However, notice that if we make a conservative choice, taking a number of variables p larger than the true one ($p > p^*$), the true relevant variables $T^*_{p^*}$ will always be included among the selected ones. This is a consequence of the fact that, by (4), (14) and $\varepsilon \in \overline{\mathcal{L}}_X^{\perp}$, the projection of Y onto $\overline{\mathcal{L}}_X$ is the same as the projection onto the span of $X(t^*_1), \ldots, X(t^*_{p^*})$. This implies that, if we are looking for the subspace of dimension $p > p^*$ that minimizes the distance of Y to its projection (see Equation (5)), this space will always include the span of $X(t^*_1), \ldots, X(t^*_{p^*})$.

Recall that we use the notation $T^* \prec T \in \Theta_p$ meaning that T^* is a sub-vector of T, that is, that the components of T^* are included within those of T. With this notation, what we have shown is that, under (14), $T^* = \left(t^*_1, \ldots, t^*_{p^*}\right) \prec \operatorname{argmax} Q_0(T_p)$, for $p \geq p^*$. In Section 4 it is addressed the problem of what to do when p^* is unknown from an empirical point of view.

2.4 A recursive expression

The function Q_0 defined in (11) can be rewritten in an alternative way, which is useful to analyze the gain when we add a new variable to a set of variables already selected. Moreover, this alternative expression paves the way for a sequential implementation of the variable selection method. To give the formula, besides the notation c_T and Σ_T, introduced earlier, we will also use c_j to denote $\operatorname{cov}(X(t_j), Y)$, σ_j^2 to denote $\operatorname{var}(X(t_j))$, and $c_{T_p,j}$ to denote the vector $(\operatorname{cov}(X(t_1), X(t_j)), \ldots, \operatorname{cov}(X(t_p), X(t_j)))'$.

Proposition 2. *Given $T_{p+1} = (t_1, \ldots, t_p, t_{p+1}) \in \Theta_{p+1}$, $p \geq 1$, and $T_p \prec T_{p+1}$, for some $p \geq 1$ such that the covariance matrices of the process $\Sigma_{T_{p+1}}$ are invertible for all $T_{p+1} \in \Theta_{p+1}$,*

$$Q_0(T_{p+1}) \;=\; Q_0(T_p) \;+\; \frac{\left(c'_{T_p} \Sigma^{-1}_{T_p} c_{T_p,p+1} - c_{p+1}\right)^2}{\sigma^2_{p+1} - c'_{T_p,p+1} \Sigma^{-1}_{T_p} c_{T_p,p+1}}. \tag{15}$$

Equation (15) is useful to simplify other derivations in the paper and it is shown in Section 8.1. Actually, it can be also proved that the quotient in (15) tends to zero when t_{p+1} tends to one of the points in T_p, so that selecting a point too close to one of those already selected is redundant and non-informative according to this criterion. The proof of this fact is quite technical so it is omitted.

Equation (15) already appears in the well-known forward selection method for variable selection in multiple regression (see e.g. Miller (2002), Section 3.2). A modification of the resulting expression is also used in the variable selection method proposed by

Yenigün and Rizzo (2015), still in the multivariate regression setting. In such alternative version, the usual covariance is replaced by the distance covariance, defined in Székely and Rizzo (2013).

The quotient in Equation (15) can be written in a more insightful way, as shown in the following result.

Proposition 3. *In the above defined setup, denoting* $X(T_p) = \big(X(t_1), \ldots, X(t_p)\big)'$, *and* $Y_{T_p} = P_{span\{X(t_i), t_i \in T_p\}} Y$ *and* $X_{T_p}(t_{p+1}) = P_{span\{X(t_i), t_i \in T_p\}} X(t_{p+1})$ *the projections on* $span\{X(t_i), t_i \in T_p\}$,

$$Q_0(T_{p+1}) \;=\; Q_0(T_p) \;+\; \frac{cov^2\big(Y - Y_{T_p},\; X(t_{p+1})\big)}{var\big(X(t_{p+1}) - X_{T_p}(t_{p+1})\big)}. \tag{16}$$

The quotient of Equation (16) is known as *part correlation coefficient* or *semi-partial correlation coefficient*, a quantity which appears in several techniques dealing with multivariate data. The proof of this result can be found in Subsection 8.2.

3 Sample properties of the variable selection method

3.1 The proposed method

In order to carry out the variable selection in practice, we have to estimate the function Q_0 from a sample $(Y_1, X_1), \ldots, (Y_n, X_n)$ of independent observations drawn from the model (4). The most natural estimator is given by $\widehat{Q}_n(T_p) = \widehat{c}'_{T_p} \widehat{\Sigma}^{-1}_{T_p} \widehat{c}_{T_p}$, where $\widehat{c}_{T_p}$ and $\widehat{\Sigma}_{T_p}$ are the sample versions of c_{T_p} and Σ_{T_p}, respectively, based on the sample mean $\overline{X}(t) = n^{-1} \sum_{i=1}^n X_i(t)$ and the sample covariances

$$\widehat{\mathrm{Cov}}(X(s), X(t)) = n^{-1} \sum_{i=1}^n X_i(s) X_i(t) - \overline{X}(s) \overline{X}(t)$$

of the trajectories. Then, if we want to select p variables we propose to use $\widehat{T}_{p,n}$, where

$$\widehat{T}_{p,n} := \underset{T_p \in \Theta_p}{\operatorname{argmax}}\; \widehat{Q}_n(T_p) = \underset{T_p \in \Theta_p}{\operatorname{argmax}}\; \widehat{c}'_{T_p} \widehat{\Sigma}^{-1}_{T_p} \widehat{c}_{T_p}. \tag{17}$$

3.2 Asymptotic results

In the following results, we will analyze the asymptotic behavior of this proposal. We start with three preliminary results that may be of some interest by themselves. First we prove that, under some moment conditions, the sample mean and covariance functions of X converge uniformly a.s. to their population counterparts:

Lemma 1. *Assume that the process* X *has continuous trajectories with continuous mean and covariance functions and that it fulfils that* $\mathbb{E}\big[\sup_{t \in [\delta,1]} X(t)^2\big] < \infty$, *for a certain* $\delta \geq 0$. *Then,*

$$\sup_{s,t \in [\delta,1]} |\widehat{Cov}(X(s), X(t)) - Cov(X(s), X(t))| \overset{a.s.}{\to} 0. \tag{18}$$

Proof. Note that the assumption implies $\mathbb{E}\big[\sup_{t\in[\delta,1]}|X(t)|\big] < \infty$ and the stochastic process $\{X(t) : t \in [\delta, 1]\}$ has finite strong expectation with trajectories in $\mathcal{C}[\delta, 1]$, which is a separable Banach space. Then, we can apply Mourier's SLLN (see e.g. Theorem 4.5.2 of Laha and Rohatgi (1979), p. 452) to conclude

$$\sup_{t\in[\delta,1]} |\overline{X}(t) - m(t)| \overset{a.s.}{\to} 0, \tag{19}$$

Similarly, the process $Z(s,t) := X(s)X(t)$, with trajectories in $\mathcal{C}([\delta,1]^2)$, is such that its strong expectation exists. Indeed, since

$$0 \leq (|X(s)| - |X(t)|)^2 = |X(s)|^2 + |X(t)|^2 - 2|Z(s,t)|,$$

it holds

$$\mathbb{E}\big(\sup_{s,t\in[\delta,1]} |Z(s,t)|\big) \leq \mathbb{E}\big(\sup_{t\in[\delta,1]} |X(t)|^2\big) < \infty.$$

Moreover, $\mathcal{C}([\delta,1]^2)$ is separable since $[\delta,1]^2$ is compact. Then, Mourier's SLLN and (19) imply (18). □

Next, we prove that both $\widehat{Q}_n$ and Q_0 are continuous functions for any $p \geq 1$:

Lemma 2. *Assume that the process $X(t)$ has continuous mean and covariance functions. Let $p \geq 1$ and $\Theta_p = \Theta_p(\delta)$ be such that the assumptions of Lemma 1 hold. In addition, asssume that the covariance matrix Σ_{T_p} is invertible for all $T_p \in \Theta_p$. Then, the functions $\widehat{Q}_n$ and Q_0 are continuous on Θ_p.*

Proof. Fix $p \geq 1$. First, we prove that Q_0 is continuous. Since the process $X(t)$ has continuous mean and covariance functions we have that

$$c_{T_p} = (\mathrm{Cov}(X(t_1), Y), \ldots, \mathrm{Cov}(X(t_p), Y))'$$

is continuous on Θ_p. On the other hand, since the entries of Σ_{T_p} are continuous on $[0,1]^2$, $\det(\Sigma_{T_p})$ is also continuous on Θ_p, where $\det(\Sigma)$ stands for the determinant of Σ. By assumption, $\det(\Sigma_{T_p}) > 0$ for all $T_p \in \Theta_p$. Since Θ_p is compact, $\inf_{T_p\in\Theta_p} \det(\Sigma_{T_p}) > 0$. Observe that

$$\Sigma_{T_p}^{-1} = \frac{\mathrm{adj}(\Sigma_{T_p})}{\det(\Sigma_{T_p})},$$

where $\mathrm{adj}(\Sigma)$ denotes the adjugate of Σ. As a consequence, the entries of $\Sigma_{T_p}^{-1}$ are continuous on Θ_p, and hence the function Q_0 is also continuous.

The proof for Q_n is analogous with the only difference that in this case we must ensure that $\inf_{T_p\in\Theta_p} \det(\widehat{\Sigma}_{T_p}) > 0$ with probability 1. Notice that for all $n \geq 1$ and $T_p \in \Theta_p$, $\det(\widehat{\Sigma}_{T_p}) > 0$ a.s. On the other hand, from (18) it follows that

$$\sup_{T_p\in\Theta_p} |\det(\widehat{\Sigma}_{T_p}) - \det(\Sigma_{T_p})| \overset{a.s}{\to} 0. \tag{20}$$

We have seen before that $\inf_{T_p\in\Theta_p} \det(\Sigma_{T_p}) > 0$. Then, with probability 1, there exists n_0 such that if $n \geq n_0$, $\inf_{T_p\in\Theta_p} \det(\widehat{\Sigma}_{T_p}) > 0$. □

The two previous lemmas allow us to prove the uniform convergence on Θ_p of the empirical criterion for variable selection to the theoretical one.

Lemma 3. *Under the assumptions of Lemma 2, it holds that*

$$\sup_{T_p \in \Theta_p} |\widehat{Q}_n(T_p) - Q_0(T_p)| \overset{a.s.}{\to} 0.$$

Proof. It is enough to establish the uniform convergence a.s. of the coordinates of $\widehat{c}_{T_p}$ and the entries of $\widehat{\Sigma}_{T_p}^{-1}$ to those of c_{T_p} and $\Sigma_{T_p}^{-1}$ respectively.

Equation (19) and the same argument leading to (19) but applied to the process $Z(t) = X(t)Y$ yield

$$\sup_{t \in [\delta, 1]} \left| \frac{1}{n} \sum_{i=1}^{n} (X_i(t) - \overline{X}(t))(Y_i - \overline{Y}) - \mathrm{Cov}(X(t), Y) \right| \overset{a.s}{\to} 0, \tag{21}$$

and hence the uniform convergence a.s. of the coordinates of $\widehat{c}_{T_p}$ to those of c_{T_p}.

Finally, observe that $\widehat{\Sigma}_{T_p}^{-1} = \frac{\mathrm{adj}(\widehat{\Sigma}_{T_p})}{\det(\widehat{\Sigma}_{T_p})}$. Then, from (18), (20) and $\inf_{T_p \in \Theta_p} \det(\Sigma_{T_p}) > 0$ we conclude the uniform convergence a.s. of the entries of $\widehat{\Sigma}_{T_p}^{-1}$ to those of $\Sigma_{T_p}^{-1}$. □

Now assume that the sparsity condition (14) holds. Then, $T_{p^*}^* = (t_1^*, \ldots, t_{p^*}^*) \in \Theta_{p^*}$ is "sufficient" in the sense that the response only depends on the regressor variable through the values $X(t_1^*), \ldots, X(t_{p*}^*)$. We have already seen that $T_{p^*}^*$ is a global maximum of Q_0 (see the remark below Equation (14)). In fact, we are going to prove that under mild conditions it is the only global maximum of Q_0 on Θ_{p^*} and that the estimator $\widehat{T}_{p^*,n}$ (defined in (17) with $p = p^*$) converges a.s. to $T_{p^*}^*$. Then, our proposal is able to identify consistently the true relevant points.

Theorem 1. *Assume* (14) *holds and that the process $X(t)$ has continuous mean and covariance functions. Suppose also that the assumptions of Lemma 1 hold for $p = p^*$, the covariance matrix $\Sigma_{T_{p^*}}$ is invertible for all $T_{p^*} \in \Theta_{p^*}$ and the covariance matrix $\Sigma_{T_{p^*} \cup S_{p^*}}$ is invertible for all $T_{p^*}, S_{p^*} \in \Theta_{p^*}$, with $T_{p^*} \neq S_{p^*}$. Then,*

(a) The point $T_{p^*}^* \in \Theta_{p^*}$, given by (14)*, is the only global maximum of Q_0 on Θ_{p^*}.*

(b) If $\widehat{T}_{p^*,n} = \mathrm{argmax}_{T_{p^*} \in \Theta_{p^*}} \widehat{Q}_n(T_{p^*})$, then $\widehat{T}_{p^*,n} \to T_{p^*}^*$ a.s. as $n \to \infty$.

(c) $\widehat{T}_{p^*,n}$ converges to $T_{p^*}^*$ in quadratic mean, that is, $\mathbb{E}\big(\|\widehat{T}_{p^*,n} - T_{p^*}^*\|_2^2\big) \to 0$, as $n \to \infty$.

Proof. (a) In view of (10), it is enough to prove that $T^* := T_{p^*}^*$ is the unique global minimum of

$$Q_1(T_{p^*}) = \|Y - Y_{T_{p^*}}\|_2^2 = \|Y_{T^*} - Y_{T_{p^*}}\|_2^2 + \mathrm{Var}(\varepsilon).$$

The expression above readily shows that T^* minimizes Q_1. Suppose that there exists another minimum $S^* \in \Theta_{p^*}$ such that $S^* \neq T^*$. Then, we must have $\|Y_{T^*} - Y_{S^*}\|_2^2 = 0$

and hence $Y_{T^*} - Y_{S^*} = 0$ a.s. As a consequence, using the notation $\widetilde{X}(t) = X(t) - m(t)$, there exist coefficients β_j and α_j such that $\sum_{j=1}^{p^*} \beta_j \widetilde{X}(t_j^*) - \sum_{j=1}^{p^*} \alpha_j \widetilde{X}(s_j^*) = 0$ a.s., for $T^*, S^* \in \Theta_{p^*}$ with $S^* \neq T^*$. This fact contradicts the assumption that the covariance matrix $\Sigma_{T^* \cup S^*}$ must be invertible. Therefore, $T^* = S^*$.

(b) Since the functions $\widehat{Q}_n$ and Q_0 are continuous on Θ_{p^*} (by Lemma 2) and the sequence of functions $\widehat{Q}_n$ tends uniformly a.s. to Q_0 on Θ_{p^*} (by Lemma 3) the fact that Q_0 has a unique maximum on Θ_{p^*}(part (a)) implies that $\widehat{T}_{p^*,n}$ converges almost surely to $T_{p^*}^*$.

(c) From part (b), we have $\|\widehat{T}_{p^*,n} - T_{p^*}^*\|_2 \to 0$ a.s. as $n \to \infty$. Moreover, since both $\widehat{T}_{p^*,n}$ and $T_{p^*}^*$ belong to Θ_{p^*},

$$\|\widehat{T}_{p^*,n} - T_{p^*}^*\|_2 \leq \|\widehat{T}_{p^*,n}\|_2 + \|T_{p^*}^*\|_2 \leq 2p^*.$$

The result follows from dominated convergence theorem (using $2p^*$ as the integrable dominating function). □

Once we have selected p^* points, we can use them to predict the response variable. The optimal predictions (in a square mean sense) are given by:

$$\widehat{Y}_{\widehat{T}_{p^*}} = \widehat{\beta}_1 \widetilde{X}(\widehat{t}_1) + \cdots + \widehat{\beta}_{p^*} \widetilde{X}(\widehat{t}_{p^*}),$$

where $\widetilde{X}(t) = X(t) - m(t)$, and $\widehat{\beta}_{\widehat{T}_{p^*}} := (\widehat{\beta}_1, \ldots, \widehat{\beta}_{p^*})' = \widehat{\Sigma}_{\widehat{T}_{p^*}}^{-1} \widehat{c}_{\widehat{T}_{p^*}}$. On the other hand, the prediction we would use under condition (14) if we knew the true relevant points and the true values of the parameters of the model would be

$$Y_{T^*} = \beta_1^* \widetilde{X}(t_1^*) + \cdots + \beta_{p^*}^* \widetilde{X}(t_{p^*}^*),$$

where now $\beta_{T_{p^*}^*}^* := (\beta_1^*, \ldots, \beta_{p^*}^*)' = \Sigma_{T_{p^*}^*}^{-1} c_{T_{p^*}^*}$. The following result refers to the asymptotic behavior of the data-driven predictions $\widehat{Y}_{\widehat{T}_{p^*}}$. It is shown that they converge a.s. and in quadratic mean to the oracle values Y_{T^*}.

Theorem 2. *Under the assumptions of Theorem 1,* $\widehat{Y}_{\widehat{T}_{p^*}} \overset{a.s.}{\to} Y_{T^*}$. *If, in addition, there exists* $\eta > 0$ *such that* $\mathbb{E}\big[\sup_{t\in[\delta,1]} |X(t)|^{2+\eta}\big] < \infty$ *then* $\widehat{Y}_{\widehat{T}_{p^*}} \xrightarrow{L^2} Y_{T^*}$, *as* $n \to \infty$.

Proof. For simplicity, denote $\widehat{T} := \widehat{T}_{p^*}$ and $T^* := T_{p^*}^*$. Observe that

$$\begin{aligned}
|\widehat{Y}_{\widehat{T}} - Y_T| =& |\widehat{c}'_{\widehat{T}} \widehat{\Sigma}_{\widehat{T}}^{-1} \widetilde{X}(\widehat{T}) - c'_{T^*} \Sigma_{T^*}^{-1} \widetilde{X}(T^*)| \\
\leq& |\widehat{c}'_{\widehat{T}} \widehat{\Sigma}_{\widehat{T}}^{-1} \widetilde{X}(\widehat{T}) - c'_{\widehat{T}} \Sigma_{\widehat{T}}^{-1} \widetilde{X}(\widehat{T})| + |c'_{\widehat{T}} \Sigma_{\widehat{T}}^{-1} \widetilde{X}(\widehat{T}) - c'_{T^*} \Sigma_{T^*}^{-1} \widetilde{X}(T^*)| \\
\leq& \|\widehat{c}'_{\widehat{T}} \widehat{\Sigma}_{\widehat{T}}^{-1} - c'_{\widehat{T}} \Sigma_{\widehat{T}}^{-1}\|_2 \, \|\widetilde{X}(\widehat{T})\|_2 + |c'_{\widehat{T}} \Sigma_{\widehat{T}}^{-1} \widetilde{X}(\widehat{T}) - c'_{T^*} \Sigma_{T^*}^{-1} \widetilde{X}(T^*)| \\
\leq& \sup_{T \in \Theta_{p^*}} \|\widehat{c}'_T \widehat{\Sigma}_T^{-1} - c'_T \Sigma_T^{-1}\|_2 \sup_{T \in \Theta_{p^*}} \|\widetilde{X}(T)\|_2
\end{aligned}$$

$$+|c'_{\widehat{T}}\Sigma^{-1}_{\widehat{T}}\widetilde{X}(\widehat{T}) - c'_{T^*}\Sigma^{-1}_{T^*}\widetilde{X}(T^*)|.$$

Then, to prove $\widehat{Y}_{\widehat{T}} \to Y_{T^*}$ a.s. it is enough to see that the two addends of the last expression goes to 0 a.s. Observe that $\sup_{T\in\Theta_{p^*}} \|\widehat{c}'_T\widehat{\Sigma}^{-1}_T - c'_T\Sigma^{-1}_T\|_2 \to 0$ a.s., as $n \to \infty$, by (18) and (21). Moreover, since $X(t)$ has continuous trajectories and continuous mean function, and Θ_{p^*} is compact, we have $\sup_{T\in\Theta_{p^*}} \|\widetilde{X}(T)\|_2 < \infty$. Finally, the continuity of c_T, Σ_T and $\widetilde{X}(T)$, together with Theorem 1(b), imply that $|c'_{\widehat{T}}\Sigma^{-1}_{\widehat{T}}\widetilde{X}(\widehat{T}) - c'_{T^*}\Sigma^{-1}_{T^*}\widetilde{X}(T^*)| \to 0$ a.s., as $n \to \infty$.

In order to prove $\widehat{Y}_{\widehat{T}} \xrightarrow{L^2} Y_{T^*}$, as $n \to \infty$, we must check that there exists $\eta > 0$ such that $\sup_n \mathbb{E}|Y_{\widehat{T}}|^{2+\eta} < \infty$, which in turn implies that the sequence $\widehat{Y}^2_{\widehat{T}}$ is uniformly integrable. Then, we can apply that a uniformly integrable sequence of random variables which converges in probability, also converges in L^1 (see e.g. Proposition 6.3.2 and the corollary of Proposition 6.3.3 in Laha and Rohatgi (1979)).

By assumption, there exists $\eta > 0$ such that $\mathbb{E}\big(\sup_{t\in[\delta,1]} |\widetilde{X}(t)|^{2+\eta}\big) < \infty$. Observe that

$$|\widehat{Y}_{\widehat{T}}|^{2+\eta} = |\widehat{c}'_{\widehat{T}}\widehat{\Sigma}^{-1}_{\widehat{T}}\widetilde{X}(\widehat{T})|^{2+\eta} \le \|\widehat{c}'_{\widehat{T}}\widehat{\Sigma}^{-1}_{\widehat{T}}\|_2^{2+\eta}\, \|\widetilde{X}(\widehat{T})\|_2^{2+\eta}.$$

We have seen that $\sup_{T\in\Theta_{p^*}} \|\widehat{c}'_T\widehat{\Sigma}^{-1}_T - c'_T\Sigma^{-1}_T\|_2 \to 0$ a.s., as $n \to \infty$. Then, given $\varepsilon > 0$, for large enough n,

$$\|\widehat{c}'_{\widehat{T}}\widehat{\Sigma}^{-1}_{\widehat{T}}\|_2^{2+\eta} \le \varepsilon + \sup_{T\in\Theta_{p^*}} \|c'_T\Sigma^{-1}_T\|_2^{2+\eta} := C < \infty.$$

From the last two displayed equations, for large enough n,

$$\mathbb{E}|\widehat{Y}_{\widehat{T}}|^{2+\eta} \le C\,\mathbb{E}\|\widetilde{X}(\widehat{T})\|\|_2^{2+\eta} \le C'\,\mathbb{E}\big(\sup_{t\in[\delta,1]} |\widetilde{X}(t)|^{2+\eta}\big) < \infty,$$

where $C' = C(p^*)^{(2+\eta)/2}$. Since the last upper bound does not depend on n, we get $\sup_n \mathbb{E}|Y_{\widehat{T}}|^{2+\eta} < \infty$. □

3.3 Sequential approach

The number of combinations of variables is usually too large to carry out an exhaustive search to find the optimal p^* variables, even for small values of p^*. Then, in practice we need to define a search strategy to perform the selection. That is, we must decide how to explore the space of all possible combinations of variables. We propose to use the sequential approach we describe in this section.

Observe that a proof analogous to that of Equations (15) and (16) also gives their corresponding sample versions:

$$\widehat{Q}_n(T_{p+1}) \;=\; \widehat{Q}_n(T_p) \;+\; \frac{\Big(\widehat{c}'_{T_p}\widehat{\Sigma}^{-1}_{T_p}\widehat{c}_{T_p,p+1} - \widehat{c}_{p+1}\Big)^2}{\widehat{\sigma}^2_{p+1} - \widehat{c}'_{T_p,p+1}\widehat{\Sigma}^{-1}_{T_p}\widehat{c}_{T_p,p+1}}, \tag{22}$$

$$\widehat{Q}_n(T_{p+1}) \;=\; \widehat{Q}_n(T_p) \;+\; \frac{\widehat{\mathrm{cov}}^2\big(Y - \widehat{Y}_{T_p},\, X(t_{p+1})\big)}{\widehat{\mathrm{var}}\big(X(t_{p+1}) - \widehat{X}_{T_p}(t_{p+1})\big)}. \tag{23}$$

These equations suggest a sequential way to carry out the variable selection. In each step, we find the variable $t_{p+1} \in [\delta, 1]$ maximizing the quotient in the equations above. In this way, we obtain nested subsets of variables, since $T_p \prec T_{p+1}$. This greedy method does not guarantee the convergence to the global maximum of $\widehat{Q}_n$, but it shows a good behavior in practice, as we will show later on.

4 Estimating the number of variables

As discussed above, our method for variable selection works, whenever the regression function β belongs to the RKHS space associated with the covariance function K. It is based on the idea of asymptotically minimizing (on $T_p = (t_1, \ldots, t_p)$) the residuals

$$Q_1(T_p) := \min_{(\beta_1,\ldots,\beta_p)\in\mathbb{R}^p} \|Y - \sum_{j=1}^{p} \beta_j \widetilde{X}(t_j)\|_2^2 = \|Y - \sum_{j=1}^{p} \beta_j^* \widetilde{X}(t_j)\|_2^2,$$

where $\widetilde{X}(t_j) = X(t_j) - m(t_j)$ and $(\beta_1^*, \ldots, \beta_d^*)' = \Sigma_T^{-1} c_T$, Σ_T being the covariance matrix of $(X(t_1), \ldots, X(t_p))'$ and c_T the vector whose j-th component is $\mathrm{cov}(X(t_j), Y)$. As proved in Proposition 1, this amounts to asymptotically maximize the function Q_0 defined in (11), which in turn is equivalent to minimize the function Q_2, defined in (9). Also, the functions, Q_1 and Q_2 agree up to an additive constant and both agree with Q_0 up to a change of sign plus an additive constant.

Throughout this section we assume the validity of the sparsity assumption (8), that is, we assume that the slope function β has the form $\beta = \sum_{j=1}^{p^*} \beta_j K(t_j^*, \cdot)$, as stated in Equation (14), for some constants $\beta_1, \ldots, \beta_{p^*} \in \mathbb{R}$ and for $T_{p^*}^* = (t_1^*, \ldots, t_{p^*}^*)$. In this case, we can properly speak of a specific target set of "true" variables $T^* = T_{p^*}^* = (t_1^*, \ldots, t_p^*)$ to be selected and, in particular, of a "true" number p^* of variables to select.

Keeping in mind these facts, the following comments provide some clues and motivation for the data-based selection of p^*. They will be formalized in the statement and proof of Lemma 4 below.

(a) On the one hand, any selection of type $T_p = (t_1, \ldots, t_p)$ with $p < p^*$ is clearly sub-optimal, since it would lack some relevant information, contributed by the variables in T^* not in T_p.

(b) Likewise, a choice T_p "by excess" with $T^* \prec T_p$ would not provide any benefit. To see this note that, under (8), the minimum of Q_2 is obviously attained at T_p^* and the value of Q_2 at such minimum is 0, which cannot be improved.

(c) As a consequence, the maximum value of Q_2 for points with $p^* + 1$ coordinates is attained at some T_{p^*+1} such that $T^* \prec T_{p^*+1}$ (that is, T^* is a sub-vector of T_{p^*+1}) but, in any case, $Q_0(T_{p^*+1}) - Q_0(T^*) = 0$.

(d) Then, the optimal p^* is such that the maximum value of $Q_0(T_p)$ agrees with that of $Q_0(T_{p^*})$ for any T_p such that $T_{p^*} \prec T_p$. Thus p^* is in fact the "elbow" value in the plot of $p \mapsto Q_0(T_p^*)$ from which on the increase of the maximum values of Q_0 stops.

The following lemma will set the theoretical basis of our procedure of estimation of p^*. As a consequence of this result, a procedure to estimate p^* is proposed in the next subsection.

Lemma 4. *Let us consider the model* (4) *under the assumption that* β *can be expressed as* $\beta(t) = \sum_{j=1}^{p} \beta_j K(t_j^*, \cdot)$*, where* p^* *is the minimal integer for which such representation holds. Define* $\widehat{Q}_n^{max}(p) = \max_{T_p \in \Theta_p} \widehat{Q}_n(T_p)$*. Then, under the assumptions of Lemma 2 we have*

(a)

$$\widehat{Q}_n^{max}(p^* + 1) - \widehat{Q}_n^{max}(p^*) \to 0, \, a.s., \tag{24}$$

where p^* *stands for the "true" number of variables in the sparse model* (14)

(b) For all $p < p^*$*,*

$$\lim_n \left(\widehat{Q}_n^{max}(p + 1) - \widehat{Q}_n^{max}(p) \right) > 0, \, a.s., \tag{25}$$

Proof. (a) Let us first prove

$$Q_0(T_{p^*+1}^*) - Q_0(T_{p^*}^*) = 0 \tag{26}$$

To see this, note that in the proof of Proposition 1 we have proved $Q_0(T_p) = \|\beta\|_K^2 - Q_2(T_p)$ with $Q_2(T_p) = \min_{(\beta_1,\ldots,\beta_p)\in\mathbb{R}^p} \|\beta - \sum_{j=1}^{p} \beta_j K(t_j, \cdot)\|_K^2$. Also, under (14), $Q_2(T_{p^*}^*) = 0$ so that $Q_0(T_{p^*}^*) = \|\beta\|_K^2$, which is the maximum possible value of Q_0. On the other hand, it is clear that $Q_2(T_{p^*+1}^*) \leq Q_2(T_{p^*}^*)$ so that we must also have $Q_2(T_{p^*+1}^*) = 0$ and $Q_0(T_{p^*+1}^*) = \|\beta\|_K^2$. This proves (26). Now conclusion (24) follows directly from the uniform convergence of Q_n to Q_0, as established in Lemma 3.

(b) Similarly to part (a) we only need to prove

$$Q_0(T_{p+1}^*) - Q_0(T_p^*) > 0 \text{ for all } p < p^*. \tag{27}$$

Indeed, assume we have $Q_0(T_{p+1}^*) - Q_0(T_p^*) = 0$ for some $p < p^*$. Then, since the prediction error $Q_1(T_p)$ defined in (5) satisfies $Q_1(T_p) = -Q_0(T_p) + \|\beta\|_K^2 + \sigma^2$ we would have that the prediction error $Q_1(T_p^*)$ obtained with p variables in the sparse model $Y = \sum_{j=1}^{q} \beta_j X(t_j) + \varepsilon$, with $\text{var}(\varepsilon) = \sigma^2$ would be the same, for $q = p$ and $q = p + 1$. Then, by recurrence, we get that the error would be in fact the same, irrespective of the number of q explanatory variables in the range $p, p + 1, \ldots, p^*$. Thus, the linear model $Y = \langle \beta, X \rangle_K + \varepsilon$ holds for a regression function of type $\beta = \sum_{j=1}^{p} \beta_j K(t_j^*, \cdot)$ with $p < p^*$. This is a contradiction with the assumption that p^* is the minimal value for which such model holds.

Now the result follows from (27) and the a.s. uniform convergence of $\widehat{Q}_n$ to Q_0. □

4.1 The estimator of p^*

The above discussion suggests the following method to estimate p^*:

1. Define
$$\Delta = \min_{p<p^*} \ (Q_0(T^*_{p+1}) - Q_0(T^*_p)) \tag{28}$$
Assume we are able to fix a value $\epsilon > 0$ such that $\epsilon < \Delta$.

2. Define
$$\widehat{p} = \min \left\{ p : \widehat{Q}^{max}_n(p+1) - \widehat{Q}^{max}_n(p) < \epsilon \right\}, \tag{29}$$

Theorem 3. *Under the assumptions of Lemma 4 the estimator $\widehat{p}$ defined in* (29) *fulfils $\widehat{p} \to p^*$, almost surely.*

Proof. This result is a direct consequence of Lemma 4 □

In practice, the calculation of $\widehat{p}$ could be made using techniques inspired in the change point detection methodology in time series. Thus, we could interpret the collection of values $L_n(p) = \log\left(\widehat{Q}^{max}_n(p+1) - \widehat{Q}^{max}_n(p)\right)$ for $p = 1, \ldots$ as a time series. Then, we could apply the usual k-means clustering algorithm to these values, with $k = 2$. We would estimate $\widehat{p}$ as the minimum value of p such that all the values $L_n(p)$ with $p \geq \widehat{p}$ belong to a different cluster than that of $L_n(1)$.

5 When p^* is not estimated: the conservative oracle property

Under the sparseness assumption (14), where p^* is unknown, another sensible approach for the choice of the number p of selected variables is to take a conservative, large enough value of p.

The basic idea of this section is easy to state: suppose that a "conservative oracle" gives us a value p such that $p > p^*$. Accordingly, we perform our variable selection procedure for such value p. This yields p variables $\widehat{t}_1, \ldots, \widehat{t}_p$. Then, we can be sure that the "true" variables $t^*_1, \ldots, t_{p^*}$ are very close to the p^* variables in $\{\widehat{t}_1, \ldots, \widehat{t}_p\}$.

The next result formalizes this property.

Theorem 4. *Let us consider the model* (4) *under the assumption that β can be expressed as $\beta(t) = \sum_{j=1}^{p^*} \beta_j K(t^*_j, \cdot)$, where p^* is the minimal integer for which such representation holds. Let $\widehat{t}_1, \ldots, \widehat{t}_p$ the variables selected by the method* (17)*, where p is a given value larger than p^*. Then, for all $\epsilon > 0$,*

$$\mathbb{P}\left(t^*_i \in \bigcup_{j=1}^{p} (\widehat{t}_j - \epsilon, \widehat{t}_j + \epsilon), \ i = 1, \ldots, p^* \right) = 1, \ \textit{eventually, as } n \to \infty. \tag{30}$$

Proof. Recall that the choice of the variables $T_p = (t_1, \ldots, t_p)$ is performed by asymptotically maximizing the function $Q_0(T_p)$, defined in (11). More precisely, as Q_0 depends on unknown population quantities, we in fact maximize the estimator $\widehat{Q}_n$ defined in Subsection 3.1.

Now, let us note that the maximum of Q_0 is not unique. Indeed, we assume that the "minimal" sparse representation of β has the form $\beta(t) = \sum_{j=1}^{p^*} \beta_j K(t_j^*, \cdot)$ but, of course, if $p > p^*$, we may formally put $\beta(t) = \sum_{i=1}^{p} \beta_i K(s_i, \cdot)$ as long as the "true" optimal points $t_i^*, \ldots, t_{p^*}$ are among the s_i's and all the coefficients β_i not matching with such t_i^*'s are null. On the other hand, from the uniqueness of the function β, all the maxima of $Q_0(T_p)$ must have an expression of this type.

Let T_p^* be one of these maxima. From the above discussion it follows that the coordinates s_i of T_p^* whose corresponding coefficients β_i are not null, must necessarily coincide with one of the points t_i^* in the minimal representation $\beta(t) = \sum_{j=1}^{p^*} \beta_j K(t_j^*, \cdot)$. Let us denote by T_p^{**} the subvector of T_p^* corresponding to these values t_i^*. Let $\widehat{T}_{p,n}^*$ be a maximizer of $\widehat{Q}_n(T_p)$ and denote by $\widehat{T}_{p,n}^{**}$ the subvector of $\widehat{T}_{p,n}^*$ corresponding to the same coordinates as those of T_p^{**}.

According to Lemma 3, $Q_n(T_p)$ converges to $Q_0(T_p)$ uniformly a.s. in T_p. This entails the a.s. convergence of the subvectors $\widehat{T}_{p,n}^{**}$ to T_p^{**}, since (as indicated above) all the maxima of Q_0 must content the subvector T_p^{**}.

Now, conclusion (30) is a direct consequence from the definition of almost sure convergence.

□

This result is reminiscent of the *Sure Screening Property* defined in Fan and Lv (2008), which is used to quantify the efficiency of multivariate variable selection methods. But, obviously, property (30) is adapted to cope with the functional nature of the data and the fact that the values t_i range on a continuous domain.

6 Experiments

The purpose of this section is to give some insights on the practical behaviour of our proposal for variable selection, both in simulations and real data examples. We are aware that the design of these experiments is largely discretionary, as the range of possible models for simulation is potentially unlimited (especially in the case of functional data models) and there is also a considerable amount of real data examples currently available in the FDA literature. Still, our choices have not been completely arbitrary. We have tried to follow some objective criteria. First, the theoretical models chosen for the simulations must obviously include some situations in which our crucial "sparseness" assumption $\beta = \sum_j \beta_j K(t_j, \cdot)$ is fulfilled. As discussed above, such models are quite natural if we are willing to use variable selection techniques. Also, it looks reasonable to include at least one model in which this assumption is not valid. Regarding the real data, we have just chosen two examples used in the recent literature for the purpose of checking other variable selection methods in functional regression settings.

In any case, we would like to emphasize that we make here no attempt to draw any definitive conclusion on the performance of our method when compared with others. In our view, no unique empirical study can lead to safe, objective conclusions in this regard: the only reliable verdict should be given by the users community, after some time of practice with real data problems. Our purposes here are far more unassuming; we just want to provide some hints suggesting that our proposal

(a) has a satisfactory performance in the "sparse" models for which it has been designed,

(b) can be implemented in practice, with an affordable computational cost,

(c) could be hopefully competitive under other theoretical models, far from the ideal assumption $\beta = \sum_j \beta_j K(t_j, \cdot)$,

(d) has also a satisfactory practical performance in a couple of real data examples commonly used in the literature of variable selection.

6.1 Simulation experiments

Keeping in mind the above general lines, we next define the simulation models under study. In our context a "model" is defined by three elements: a stochastic process (from which the functional data are generated), a regression equation, of type $Y = \langle X, \beta \rangle_K + \varepsilon$ (or, more generally, $Y = g(X) + \varepsilon$) and an error variable ε. In what follows, ε has been chosen in all cases as $\varepsilon \sim N(0, \sigma)$ with $\sigma = 0.2$.

We have considered six processes, covering a broad range of different situations.

1. *Standard Brownian Motion* (Bm) $\{B(t),\ t \in [0,1]\}$.

2. *Geometric Brownian Motion* (gBm). This non-Gaussian process is also known as exponential Brownian motion. It can be defined just by $X(t) = e^{B(t)}$.

3. *Integrated Brownian Motion* (iBm): it is obtained as $X(t) = \int_0^t B(s)\mathrm{d}s$. Note that the trajectories of this non-Markovian process are smooth.

4. *Ornstein-Uhlenbeck process* (OU). This is a Gaussian process $\{X(t)\}$ which satisfies the stochastic differential equation $\mathrm{d}X(t) = \theta(\mu - X(t))\mathrm{d}t + \sigma \mathrm{d}B(t)$. In our simulations we have chosen $\theta = \mu = \sigma = 1$.

5. *Fractional Brownian Motion* (fBM). This process is an generalization of the Brownian motion $B(t)$ but, unlike $B(t)$, it has not independent increments. The mean function of this Gaussian process is identically 0 and its covariance function is $k(t,s) = \frac{1}{2}(|t|^{2H} + |s|^{2H} - |t-s|^{2H})$, where $H \in (0,1)$ is the so-called Hurst exponent. Note that for $H = 0.5$, this process coincides with the standard Brownian Motion. Also, the trajectories of this process are still not differentiable at every point but the index H is closely related to the Hölder continuity properties of these trajectories. This entails when $H > 0.5$, the trajectories look "more regular" than those of the Brownian motion, having a wilder appearance for $H < 0.5$. To cover both cases we have used $H = 0.2$ and $H = 0.8$ in our simulations.

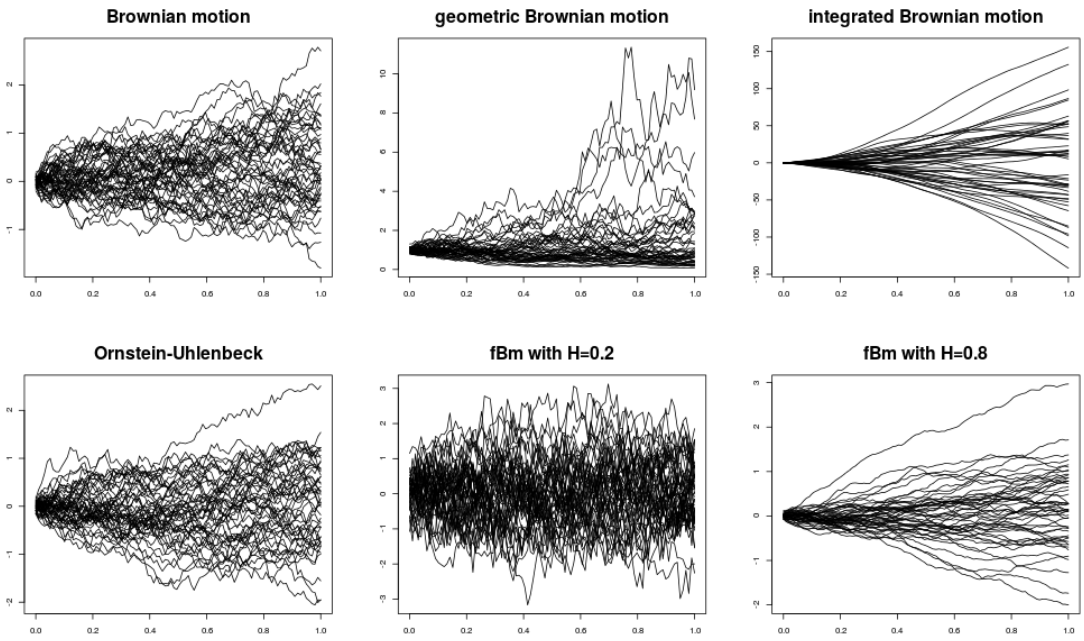

Figure 1: 50 trajectories of each of the processes used in the simulations.

Figure 1 shows some trajectories of each of these six processes, where the variables $X(t)$ for t in a neighbourhood of 0 have been omitted to satisfy the non-degeneracy requirements of the method.

As for the regression function g we have considered the following three choices.

1. Two functions β in (4) of type $\beta(t) = \sum_j \beta_j K(t_j^*, \cdot)$ so that the regression model reduces to the "sparse version" (8). We have considered two different regression functions. For the first one, we have used the set of points $T^* = (0.2, 0.4, 0.9)$ with weights $\beta_{T^*} = (2, -5, 1)$; this is "Regression model 1" in the tables. For the second one, we have used $T^* = (0.16, 0.47, 0.6, 0.85, 0.91)$ and $\beta_{T^*} = (2.1, -0.2, -1.9, 5, 4.2)$ ("Regression model 2" in the tables). Therefore, the response variables in both cases are given, respectively, by

$$\begin{aligned} Y_1 &= 2X(0.2) - 5X(0.4) + X(0.9) + \varepsilon, \\ Y_2 &= 2.1X(0.16) - 0.2X(0.47) - 1.9X(0.67) + 5X(0.85) + 4.2X(0.91) + \varepsilon. \end{aligned}$$

2. $\beta(t) = \log(1 + t)$ and the regression model is (1) with $\alpha_0 = 0$. Thus, the sparse RKHS model (8) does not hold in this case. This is "Regression model 3" in the tables. It has been already used in Cuevas et al. (2002). Therefore, the corresponding response variable is generated by

$$Y_3 = \int_0^1 \log(1 + t)X(t)\mathrm{d}t + \varepsilon.$$

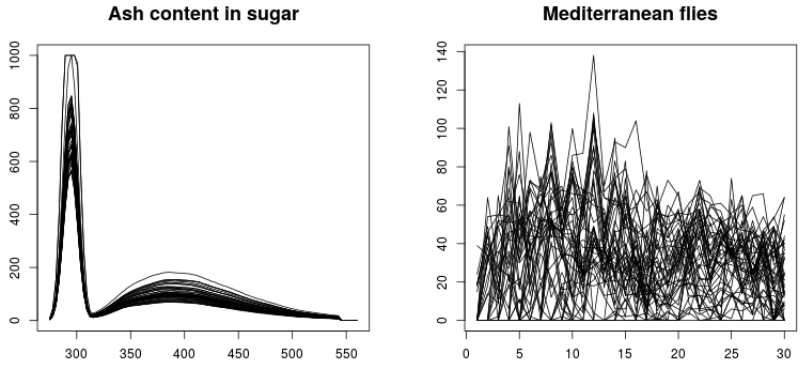

Figure 2: 50 trajectories of each of the real data sets.

6.2 Real data

We have also checked the different methods when applied to two real data sets. Since these data have been already considered in other recent papers of the FDA literature, we will give only brief descriptions of them.

1. *Ash content in sugar samples.* This data set has been used, for example, in Aneiros and Vieu (2014). The version we use corresponds in fact to a subset of the whole data set, available in `http://www.models.kvl.dk/Sugar_Process`. The response variable Y is the percentage of ash content in 266 sugar samples. The trajectories $X(t)$ are the fluorescence spectra from 275 to 560 nm at excitation wavelength 290. These curves are discretized in a grid of 100 equispaced points.

2. *Mediterranean fruit flies.* This data set has been recently used, for example, in Ji and Müller (2016). The trajectories $X(t)$ are the egg-laying profiles (number of eggs at the day t) for 512 female Mediterranean fruit flies during a period of 30 days, and the response variable Y is the total number of eggs laid during their remaining lifetime. Therefore, the trajectories are discretized in a grid of 30 points.

Figure 2 shows some trajectories of both data sets.

6.3 Methods under study and methodology

We compare our proposal with other methods for variable selection recently considered in the literature which, in general, have been implemented in the original papers. We now list the methods under study along with the notation used in the tables below.

1. The method proposed in this paper (RKHS). It has been implemented using the iterative approximation described in equation (23). The number of relevant points is chosen as explained in Section 4. Therefore, no validation technique is required.

2. The variable selection procedure proposed in Kneip et al. (2016) (KPS): in the original article, a mixed method for standard functional linear regression and variable selection technique is proposed. Since we are here concerned with variable selection, we have implemented just the corresponding part of the proposal. Essentially, the idea is to select the points (called "impact points" in Kneip et al. (2016)) maximizing the covariance between the response variable Y and a "decorrelated" version, $Z(t)$, of the original process. By construction, the decorrelated process $Z(t)$ is such that $Z(t)$ and $Z(s)$ are almost uncorrelated whenever $|t-s| \geq \delta$. The value δ and the number of selected variables are chosen using the BIC criterion, as proposed in the original paper.

3. *Partitioning Variable Selection* (PVS) with ML penalization, as proposed in Aneiros and Vieu (2014). The original sample must be split into two independent subsamples, which should be asymptotically of the same sizes. The basic idea is to apply some multivariate variable selection technique in this context, but taking advantage of the functional structure of the data. The procedure works in two steps. In the first step, one constructs a equispaced subgrid of variables among all the variables in the original grid. Then a variable selection technique for multivariate data is applied on this subgrid, using the first subsample of the original data. For instance, we might use LASSO with ML penalization, as proposed in the original paper. Then, in the second step, this variable selection technique is applied again to an enlarged grid, constructed by taking all variables in an interval around those selected in step 1 (using the second subsample). Since the version of LASSO that we are considering selects automatically the number of variables, the only smoothing parameter to be selected here is the grid step for the points used in first stage of the method. This parameter is set by cross-validation.

4. *Maxima Hunting* (MH), proposed in Berrendero et al. (2016): The original method was proposed for variable selection in supervised classification, but there is no conceptual restriction to apply the same procedure in a regression setting. The basic idea is to select the local maxima of the "distance covariance" (association measure for random variables introduced in Székely et al. (2007)) between the response and the marginal variables of the process. In practice, the numerical identification of these maxima depends on a smoothing parameter h which is chosen by cross-validation. The number of variables is also set by cross-validation.

5. *Partial Least Squares* (PLS). This technique is well-known among the FDA users of functional data. The goal of PLS is not to pick up a few variables but to select some appropriate linear functionals of the original data (very much in the spirit of principal components analysis). So PLS is not a variable selection procedure, but a dimension reduction method. This means that PLS is not directly comparable to the variable selection methods considered here, since its aims are not the exactly same. When we choose to use a variable selection procedure, it is understood that we want to perform some kind of dimension reduction still keeping the interpretability of the information directly given in terms of the original variables. By

contrast, PLS might perhaps provide some gains in efficiency but at the expense of doing a dimensionality reduction with a more difficult interpretation. Anyway, we have included PLS in our study as a useful reference, just in order to check how much do we lose by restricting ourselves to variable selection methods. We have used the function "fregre.pls.cv" of the fda.usc R-package to compute the predictions.

6. *Base.* We denote by "base" the prediction methodolgy derived from the standard L_2 linear regression model (1). No variable selection or dimension reduction procedure is done. Hence, this method is incorporated just for the sake of comparison, to assess the accuracy of the predictions based on some previous variable selection procedure with those provided by the standard functional regression model (1). We have used the function "fregre.basis.cv" of the fda.usc R-package (which relies on a basis representation of the trajectories) to compute the predictions.

For each model, all methods are checked for the sample size $n = 150$, which has been split on 100 observations used as training data and 50 employed as test data. As usual, the functional simulated data are discretized to $(X(t_1), \ldots, x(t_{100}))$, where t_i are equispaced points in $[0, 1]$, starting from $t_1 = 1/100$. For all the experiments we obtain the Relative Mean Squared Error (RMSE) of each method, as defined by

$$RMSE\left(\widehat{Y}; Y\right) = \frac{\sum_{i=1}^{n}\left(\widehat{Y}_i - Y_i\right)^2}{\sum_{i=1}^{n}|Y_i|}. \tag{31}$$

Moreover, in those cases where β has a "sparse" form of type $\beta(t) = \sum_j \beta_j K(t_j, \cdot)$ we obtain two measures of the accuracy in the variable selection procedure. Namely, we calculate the Hausdorff distance between the set of estimated points and the set of the real ones T^* as well as the number of points selected $(\widehat{p})$ in order to compare it with the real number (p^*). We have imposed a maximum of 10 selected points to the methods, except to PVS, since this method does not permit to decide the number of selected variables. The Hausdorff distance gives us an idea of the precision of the method since it takes into account the separation to the real points as well as the number of selected points. Each experiment has been replicated 100 times.

6.4 Numerical outputs

Tables 1 and 2 provide the performance of the methods measured in terms of the Relative Mean Squared Error (Eq. (31)) of the predictions, for each of the three regression models tested and for the two real data sets respectively. Methods are presented in rows. Models appear in columns. Table 3 contents the Hausdorff distance between the selected and the "true" relevant points for the four variable selection methods and the two models with sparse β function. For these same experiments, Table 4 provides the number of selected points. In all the tables of the simulated data appear the mean and the standard deviation of each of the measured quantities.

Regression model 1

	Bm	gBm	iBm	OU	fBm 0.2	fBm 0.8
RKHS	0.313 (0.0813)	0.183 (0.0885)	0.0498 (0.0409)	0.304 (0.0827)	0.185 (0.0479)	0.501 (0.127)
KPS	1.02 (0.683)	3.9 (3.71)	0.682 (0.809)	1.01 (0.823)	1.38 (1.41)	0.645 (0.185)
PVS	0.354 (0.0869)	0.369 (0.6430)	0.587 (0.218)	0.358 (0.08660)	0.235 (0.069)	0.551 (0.15)
MH	1.64 (0.841)	3.29 (3.56)	227 (47.2)	1.58 (0.839)	6.23 (1.63)	3.18 (1.31)
PLS	1.57 (0.473)	2.36 (0.981)	0.618 (0.207)	1.53 (0.438)	9.61 (2.59)	0.736 (0.181)
Base	1.88 (0.419)	2.59 (1.19)	0.134 (0.0298)	1.85 (0.464)	13.3 (3.02)	0.693 (0.175)

Regression model 2

	Bm	gBm	iBm	OU	fBm 0.2	fBm 0.8
RKHS	0.0909 (0.0257)	0.0425 (0.013)	0.00438 (0.00468)	0.0876 (0.0202)	0.0983 (0.0226)	0.0960 (0.0241)
KPS	1.63 (0.863)	3.23 (2.14)	0.14 (0.206)	1.65 (0.794)	6.83 (4.77)	0.211 (0.127)
PVS	0.179 (0.0626)	0.448 (0.488)	5.96 (1.71)	0.171 (0.0686)	0.202 (0.136)	0.206 (0.111)
MH	2.63 (0.9)	5.78 (4.98)	2.88 (0.644)	2.36 (0.979)	7.27 (5.07)	0.429 (0.14)
PLS	0.767 (0.210)	1.48 (0.838)	0.227 (0.0737)	0.757 (0.202)	6.94 (1.78)	0.151 (0.0385)
Base	0.731 (0.152)	1.63 (1.03)	0.0122 (0.00307)	0.741 (0.207)	9.01 (2.25)	0.125 (0.0272)

Regression model 3

	Bm	gBm	iBm	OU	fBm 0.2	fBm 0.8
RKHS	2.26 (0.458)	1.37 (0.489)	0.0621 (0.0144)	2.20 (0.480)	2.15 (0.46)	2.35 (0.512)
KPS	2 (0.384)	1.03 (0.292)	0.0613 (0.0202)	1.91 (0.345)	1.99 (0.405)	2.03 (0.376)
PVS	1.91 (0.389)	1.1 (0.36)	0.215 (0.0573)	1.88 (0.357)	1.94 (0.417)	2 (0.421)
MH	1.83 (0.294)	1.18 (0.551)	0.706 (0.151)	1.79 (0.295)	1.9 (0.421)	1.9 (0.358)
PLS	1.75 (0.291)	0.847 (0.221)	0.0506 (0.0111)	1.72 (0.3)	1.8 (0.407)	1.82 (0.337)
Base	1.82 (0.345)	0.885 (0.257)	0.052 (0.0121)	1.78 (0.294)	1.72 (0.373)	1.98 (0.409)

Table 1: Mean and standard deviation of the RMSE for the response variable for simulated data sets (scale of 10^{-4}).

In view of these tables, for the simulated data the proposed method seems to outperform the other variable selection procedures, according to all the considered criteria metrics (RMSE, Hausdorff distance and number of points) whenever a sparse model of type (2) holds. The PVS method also performs quite well. When the model is not satisfied, and the response variable depends on the whole trajectory, PLS is the best in general, as expected. However, the order of magnitude of the error for PLS is the same as that of the variable selection methods in most cases. That is, using a few number of variables instead of the whole trajectory does not significantly affect the prediction error, even if the response depends on the whole trajectory. For the real data, RKHS and KPS are the winners for the sugar data set. For the Mediterranean fruit flies set, the errors of all the variable selection techniques are very similar to the ones obtained with PLS and the base method.

On the other hand, regarding with the execution time, we also provide a couple of results. We have measured the execution time of the six methods for the third regression model when the functional data are drawn from the Ornstein-Uhlenbeck process and the fractional Brownian motion with $H = 0.8$. The results can be seen in Table 5. As we have already mentioned, the RKHS-based method does not require any validation step to determine the number of selected variables. Therefore, the execution time is significantly smaller than that of the other variable selection methods. We can also see that the execution time for the PVS method is much bigger than the others. Note

	Ash content in sugar	Mediterranean fruit flies
RKHS	0.339	0.945
KPS	0.375	0.967
PVS	0.616	0.996
MH	0.812	0.892
PLS	0.662	0.977
Base	0.412	0.946

Table 2: RMSE for the response variable for real data sets.

Regression model 1

	Bm	gBm	iBm	OU	fBm 0.2	fBm 0.8
RKHS	0.105 (0.15)	0.123 (0.0194)	0.991 (0.0514)	0.097 (0.162)	0 (0)	0.203 (0.208)
KPS	2.15 (0.35)	2.18 (0.217)	2.15 (0.281)	2.12 (0.285)	2.11 (0.4)	2.14 (0.277)
PVS	1.67 (0.539)	1.65 (0.494)	0.903 (0.619)	1.71 (0.411)	1.73 (0.489)	1.62 (0.537)
MH	2.05 (0.44)	2.07 (0.517)	5.91 (0.849)	2.10 (0.623)	2.01 (0.362)	3.03 (1.72)

Regression model 2

	Bm	gBm	iBm	OU	fBm 0.2	fBm 0.8
RKHS	1.29 (0.0294)	1.24 (0.167)	0.672 (0.316)	2.03 (0.0911)	1.3 (0)	1.26 (0.182)
KPS	1.29 (0.195)	1.68 (1.2)	1.28 (0.175)	2.19 (0.22)	1.28 (0.24)	1.27 (0.151)
PVS	0.977 (0.369)	1.18 (4.48)	5.79 (2.52)	2.03 (0.0611)	1.02 (0.408)	1.42 (0.834)
MH	1.53 (0.826)	1.43 (1.43)	7.22 (0.0851)	2.14 (0.386)	1.38 (0.428)	2.45 (0.943)

Table 3: Mean and standard deviation of the Hausdorff distance to the actual relevant points (Scale of 10^{-1}).

Regression model 1 ($p^* = 3$)

	Bm	gBm	iBm	OU	fBm 0.2	fBm 0.8
RKHS	3.21 (0.409)	3.63 (0.774)	6.27 (0.468)	3.22 (0.426)	3 (0)	3.32 (0.529)
KPS	8.94 (1.6)	9.2 (1.25)	9.66 (0.781)	8.64 (1.77)	8.55 (1.85)	9.21 (1.31)
PVS	11.42 (4.163)	10.96 (4.1)	11.61 (3.39)	12.13 (4.23)	11.36 (4.3)	11.1 (4.29)
MH	5.95 (1.546)	6.23 (1.78)	1.15 (0.359)	5.84 (1.61)	5.92 (1.61)	2.88 (1.16)

Regression model 2 ($p^* = 5$)

	Bm	gBm	iBm	OU	fBm 0.2	fBm 0.8
RKHS	4.91 (0.771)	5.5 (1.22)	6.81 (1.09)	5.06 (0.874)	4 (0)	5.23 (0.617)
KPS	9.7 (0.706)	9.46 (1.16)	9.58 (0.878)	9.48 (0.979)	9.33 (1.49)	9.66 (0.901)
PVS	8.88 (2.16)	8.77 (2.5)	7.73 (3.74)	9.23 (1.76)	7.28 (2.02)	10.61 (3.14)
MH	5.7 (1.85)	6.02 (1.73)	1.08 (0.273)	5.96 (1.73)	6.61 (2.22)	3.53 (1.11)

Table 4: Mean and standard deviation of the number of selected points ($\widehat{p}$).

	OU	fBm 0.8
RKHS	$\boxed{0.224}$ (0.051)	$\boxed{0.244}$ (0.060)
KPS	3.054 (9.637)	2.738 (0.319)
PVS	36.622 (12.462)	23.531 (4.829)
MH	1,247 (0.055)	1.299 (0.083)
PLS	0.239 (0.062)	0.246 (0.058)
Base	$\boxed{0.228}$ (0.051)	$\boxed{0.235}$ (0.055)

Table 5: Mean and standard deviation of the execution time.

however that this method has in general a good behaviour in terms of prediction error.

7 Conclusions

The RKHS approach we have introduced in this paper provides a natural framework for a formal unified theory of variable selection for functional data. The "sparse" models (those where the variable selection techniques are fully justified) appear as particular cases in this setup. As a consequence, it is possible to derive asymptotic consistency results as those obtained in the paper. Likewise, it is also possible to consider the problem of estimating the "true" number of relevant variables in a consistent way, as we do in Section 5. This is in contrast with other standard proposals for which the number of variables is previously fixed as an input, or it is determined using cross validation and other computationally expensive methods. Then, our proposal is more firmly founded in theory and, at the same time, provides a much faster method in practice, which is important when dealing with large data sets.

The empirical results we have obtained are encouraging. In short, according to our experiments, the RKHS-based method works better than other variable selection methods is those sparse models that fulfil the ideal theoretical conditions we need. In the non sparse model considered in the simulations, the RKHS method is slightly outperformed by other proposals (but still behaves reasonably). Finally, in the "neutral" field of real data examples the performance looks also satisfactory and competitive.

Last but not least, from a general, methodological point of view, this paper represents an additional example of the surprising usefulness of reproducing kernels in statistics. Additional examples can be found in Berlinet and Thomas-Agnan (2004), Hsing and Eubank (2015), Berrendero et al. (2015) or Kadri et al. (2015).

8 Some additional proofs

8.1 Proposition 2

We have to rewrite the expression $c'_{T_{p+1}}\Sigma^{-1}_{T_{p+1}}c_{T_{p+1}}$, where $p \geq 1$. We can write the matrix $\Sigma_{T_{p+1}}$ in block form as

$$\Sigma_{T_{p+1}} = \left(\begin{array}{ccc|c} & \Sigma_{T_p} & & \begin{array}{c}\mathrm{cov}(X(t_1),X(t_{p+1}))\\ \vdots \\ \mathrm{cov}(X(t_p),X(t_{p+1}))\end{array} \\ \hline \mathrm{cov}(X(t_1),X(t_{p+1})) & \ldots & \mathrm{cov}(X(t_p),X(t_{p+1})) & \mathrm{cov}(X(t_{p+1}),X(t_{p+1})) \end{array}\right)$$

$$\equiv \left(\begin{array}{c|c} \Sigma_{T_p} & c_{T_p,p+1} \\ \hline c'_{T_p,p+1} & \sigma^2_{p+1} \end{array}\right).$$

Then its inverse matrix is

$$\Sigma^{-1}_{T_{p+1}} = \left(\begin{array}{c|c} \Sigma^{-1}_{T_p} + \frac{1}{a}\Sigma^{-1}_{T_p}c_{T_p,p+1}c'_{T_p,p+1}\Sigma^{-1}_{T_p} & -\frac{1}{a}\Sigma^{-1}_{T_p}c_{T_p,p+1} \\ \hline -\frac{1}{a}c'_{T_p,p+1}\Sigma^{-1}_{T_p} & \frac{1}{a} \end{array}\right),$$

where $a = \sigma^2_{p+1} - c'_{T_p,p+1}\Sigma^{-1}_{T_p}c_{T_p,p+1}$. We can also write the vector of covariances as

$$c'_{T_{p+1}} = (\mathrm{cov}(X(t_1),Y),\ldots,\mathrm{cov}(X(t_{p+1}),Y)) = (c_{T_p} \mid c_{p+1}).$$

Using this notation we can rewrite the original expression as follows,

$$\begin{aligned} c'_{T_{p+1}}\Sigma^{-1}_{T_{p+1}}c_{T_{p+1}} &= c'_{T_p}\Sigma^{-1}_{T_p}c_{T_p} + \frac{1}{a}c'_{T_p}\Sigma^{-1}_{T_p}c_{T_p,p+1}c'_{T_p,p+1}\Sigma^{-1}_{T_p}c_{T_p} - \frac{c_{p+1}}{a}c'_{T_p,p+1}\Sigma^{-1}_{T_p}c_{T_p} \\ &\quad -\frac{c_{p+1}}{a}c'_{T_p}\Sigma^{-1}_{T_p}c_{T_p,p+1} + \frac{c^2_{p+1}}{a} \\ &= c'_{T_p}\Sigma^{-1}_{T_p}c_{T_p} + \frac{1}{a}\left[\left(c'_{T_p}\Sigma^{-1}_{T_p}c_{T_p,p+1}\right)^2 - 2c_{p+1}c'_{T_p}\Sigma^{-1}_{T_p}c_{T_p,p+1} + c^2_{p+1}\right] \\ &= c'_{T_p}\Sigma^{-1}_{T_p}c_{T_p} + \frac{\left(c'_{T_p}\Sigma^{-1}_{T_p}c_{T_p,p+1} - c_{p+1}\right)^2}{\sigma^2_{p+1} - c'_{T_p,p+1}\Sigma^{-1}_{T_p}c_{T_p,p+1}}, \end{aligned}$$

since the product $c'_{T_p}\Sigma^{-1}_{T_p}c_{T_p,p+1}$ is actually a real number.

8.2 Proposition 3

Using the notation of the statement, if $\widetilde{X}(t)$ is the centred process, we can rewrite the numerator of the quotient of Equation (15) as

$$\begin{aligned}\operatorname{cov}\big(Y - Y_{T_p}, X(t_{p+1})\big) &= c_{p+1} - \operatorname{cov}\big(Y_{T_p}, X(t_{p+1})\big) \\ &= c_{p+1} - c'_{T_p}\Sigma_{T_p}^{-1}\operatorname{cov}\big(\widetilde{X}(T_p), X(t_{p+1})\big) \\ &= c_{p+1} - c'_{T_p}\Sigma_{T_p}^{-1}c_{T_p,p+1},\end{aligned}$$

since the covariances are not affected by the centring. We can also write $\operatorname{cov}\big(Y - Y_{T_p}, X(t_{p+1}) - X_{T_p}(t_{p+1})\big)$ in the same way, since $(Y - Y_{T_p}) \perp X_{T_p}(t_{p+1})$. For the denominator,

$$\begin{aligned}&\operatorname{var}\big(X(t_{p+1}) - X_{T_p}(t_{p+1})\big) \\ &= \operatorname{var}\big(X(t_{p+1})\big) + \operatorname{var}\big(X_{T_p}(t_{p+1})\big) - 2\operatorname{cov}(X(t_{p+1}), X_{T_p}(t_{p+1})) \\ &= \sigma^2_{p+1} + c'_{T_p,p+1}\Sigma_T^{-1}c_{T_p,p+1} - 2c'_{T_p,p+1}\Sigma_T^{-1}\operatorname{cov}\left(X(t_{p+1}), \widetilde{X}(T_p)\right) \\ &= \sigma^2_{p+1} - c'_{T_p,p+1}\Sigma_T^{-1}c_{T_p,p+1}.\end{aligned}$$

From these two expressions the proposition follows straightforwardly.

Acknowledgements

This work has been partially supported by Spanish Grant MTM2016-78751-P and the European Social Fund (Ayudas para contratos predoctorales para la formación de doctores 2015, Ministerio de Economía y Competitividad, Spain).

References

G. Aneiros and P. Vieu. Variable selection in infinite-dimensional problems. *Statistics and Probability Letters*, 94:12–20, 2014.

A. Berlinet and C. Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Kluwer Academic, Boston, 2004. ISBN 1-4020-7679-7. URL `http://opac.inria.fr/record=b1128469`.

J. R. Berrendero, A. Cuevas, and J. L. Torrecilla. On near perfect classification and functional fisher rules via reproducing kernels. *arXiv preprint arXiv:1507.04398*, 2015.

J. R. Berrendero, A. Cuevas, and J. L. Torrecilla. Variable selection in functional data classification: a maxima-hunting proposal. *Statistica Sinica*, 26:619–638, 2016.

H. Cardot and P. Sarda. Functional linear regression. In F. Ferraty and Y. Romain, editors, *Handbook of Functional Data Analysis*, pages 21–46. Oxford University Press, Oxford, 2010.

A. Cuevas. A partial overview of the theory of statistics with functional data. *Journal of Statistical Planning and Inference*, 147:1–23, 2014.

A. Cuevas, M. Febrero, and R. Fraiman. Linear functional regression: The case of fixed design and functional response. *The Canadian Journal of Statistics / La Revue Canadienne de Statistique*, 30(2):285–300, 2002. ISSN 03195724. URL `http://www.jstor.org/stable/3315952`.

J. Fan and J. Lv. Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(5): 849–911, 2008.

J. Fan and J. Lv. A selective overview of variable selection in high dimensional feature space. *Statistica Sinica*, 20(1):101, 2010.

F. Ferraty and P. Vieu. *Nonparametric functional data analysis: theory and practice*. Springer Science & Business Media, 2006.

L. Horváth and P. Kokoszka. *Inference for functional data with applications*, volume 200. Springer Science & Business Media, 2012.

T. Hsing and R. Eubank. *Theoretical foundations of functional data analysis, with an introduction to linear operators*. John Wiley & Sons, 2015.

S. Janson. *Gaussian hilbert spaces*, volume 129. Cambridge university press, 1997.

H. Ji and H.-G. Müller. Optimal designs for longitudinal and functional data. *J. Roy. Statist. Soc., B, to appear*, 2016.

H. Kadri, E. Duflos, P. Preux, S. Canu, A. Rakotomamonjy, and J. Audiffren. Operator-valued kernels for learning from functional response data. *Journal of Machine Learning Research*, 16:1–54, 2015.

A. Kneip, D. Poss, and P. Sarda. Funcional linear regression with points of impact. *Ann. Statist.*, 44(1):1–30, 2016.

R. Laha and V. Rohatgi. *Probability Theory*. John Wiley & Sons, New York-Chichester-Brisbane, Wiley Series in Probability and Mathematical Statistics, 1979.

M. N. Lukić and J. H. Beder. Stochastic processes with sample paths in reproducing kernel hilbert spaces. *Transactions of the American Mathematical Society*, 353(10):pp. 3945–3969, 2001. ISSN 00029947. URL `http://www.jstor.org/stable/2693779`.

I. W. McKeague and B. Sen. Fractals with point impact in functional linear regression. *Ann. Statist.*, 38(4):2559–2586, 08 2010. doi: 10.1214/10-AOS791. URL `http://dx.doi.org/10.1214/10-AOS791`.

A. J. Miller. *Subset Selection in Regression*. Monographs on statistics and applied probability. Chapman & Hall/CRC, Boca Raton, 2002. ISBN 1-58488-171-2. URL `http://opac.inria.fr/record=b1131181`.

E. Parzen. Regression analysis of continuous parameter time series. In *Proceedings of the Fourth Berkeley Symposion on Mathematical Statistics and Probability*, volume 1, pages 469–489, 1961.

N. S. Pillai, Q. Wu, F. Liang, S. Mukherjee, and R. L. Wolpert. Characterizing the function space for bayesian kernel models. *Journal of Machine Learning Research*, 8 (Aug):1769–1797, 2007.

G. Székely and M. Rizzo. Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, 143:1249–1272, 2013.

G. J. Székely, M. L. Rizzo, and N. K. Bakirov. Measuring and testing dependence by correlation of distances. *Ann. Statist.*, 35(6):2769–2794, 12 2007. doi: 10.1214/009053607000000505. URL `http://dx.doi.org/10.1214/009053607000000505`.

C. D. Yenigün and M. L. Rizzo. Variable selection in regression using maximal correlation and distance correlation. *Journal of Statistical Computation and Simulation*, 85(8): 1692–1705, 2015.